arXiv:1601.03240v2 [cs.DB] 20 Apr 2016

Counting Answers to Existential Positive Queries: A Complexity Classification

Hubie Chen
Universidad del País Vasco, E-20018 San Sebastián, Spain
and IKERBASQUE, Basque Foundation for Science, E-48011 Bilbao, Spain

Stefan Mengel
CNRS, CRIL UMR 8188, France

Abstract

Existential positive formulas form a fragment of first-order logic that includes and is semantically equivalent to unions of conjunctive queries, one of the most important and well-studied classes of queries in database theory. We consider the complexity of counting the number of answers to existential positive formulas on finite structures and give a trichotomy theorem on query classes, in the setting of bounded arity. This theorem generalizes and unifies several known results on the complexity of conjunctive queries and unions of conjunctive queries.

1 Introduction

1.1 Background

The computational problem of evaluating a formula (of some logic) on a finite relational structure is of central interest in database theory and logic. In the context of database theory, this problem is often referred to as *query evaluation*, as it models the posing of a query to a database, in a well-acknowledged way: the formula is the query, and the structure represents the database. We will refer to the results of such an evaluation as *answers*; logically, these are the satisfying assignments of the formula on the structure. The particular case of this problem where the formula is a sentence is often referred to as *model checking*, and even in just the case of first-order sentences, can capture a variety of well-known decision problems from all throughout computer science [FG06].

In this article, we study the counting version of this problem, namely, given a formula and a structure, output the *number* of answers (see for example [PS11, GS14, DM13, CM14a] for previous studies). This problem of *counting query answers* generalizes model checking, which can be viewed as the particular case thereof where one is given a sentence and structure, and wants to decide if the

number of answers is 1 or 0, corresponding to whether or not the empty assignment is satisfying. In addition to the counting problem's basic and fundamental interest, it can be pointed out that all practical query languages supported by database management systems have a counting operator. Indeed, it has been argued that database queries with counting are at the basis of decision support systems that handle large data volume [GS14].

As has been previously articulated, a typical situation in the database setting is the evaluation of a relatively short formula on a relatively large structure. Consequently, it has been argued that, in measuring the time complexity of query evaluation tasks, one could reasonably allow a slow (non-polynomial-time) preprocessing of the formula, so long as the desired evaluation can be performed in polynomial time following the preprocessing [PY99, FG06]. Relaxing polynomial-time computation to allow arbitrary preprocessing of a *parameter* of a problem instance yields, in essence, the notion of *fixed-parameter tractability*. This notion of tractability is at the core of *parameterized complexity theory*, which provides a taxonomy for classifying problems where each instance has an associated parameter. We make use of this paradigm in this article; here, the formula is the parameter.

1.2 Contribution

Existential positive queries are the first-order formulas built from the two binary connectives $(\wedge, \vee)$ and existential quantification. They include and are semantically equivalent to the so-called *unions of conjunctive queries*, also known as *select-project-join-union queries*; these have been argued to be the most common database queries [AHV95]. Indeed, each union of conjunctive queries can be viewed as an existential positive query having a particular form, namely, a disjunction of primitive positive formulas; recall that a *primitive positive* query is an existential positive query that does not use disjunction.

We study the problem of counting query answers on existential positive queries. An established way to understand which types of queries are computationally well-behaved and exhibit desirable, tractable behavior is to consider this problem relative to a set of queries, and to attempt to understand on which sets this problem is tractable. Precisely, each set Φ of existential positive queries yields a restricted version of the general problem, namely: count the number of answers of a given formula $\phi \in \Phi$ on a given finite structure $\mathbf{B}$. We hence have a family of problems, one problem for each such set Φ. Our study focuses on formula sets that have *bounded arity* (by which is meant that there is a constant that upper bounds the arity of all relation symbols used in formulas); let us assume this property of all formula sets in this discussion.[1]

In this article, we prove a trichotomy theorem (Theorem 3.2) on the parameterized complexity of the discussed family of problems, which describes the complexity of every such problem. In particular, our trichotomy theorem

[1] Note that in the case of unbounded arity, complexity may depend on the choice of representation of relations [CG10].

shows that—in a sense made precise—each such problem is fixed-parameter tractable, equivalent to the *clique* problem, or as hard as the *counting clique* problem (which generalizes the clique problem). Note that the hypothesis that the clique problem is not fixed-parameter tractable is an established one in parameterized complexity;[2] under this hypothesis, our trichotomy theorem yields a precise description of the problems (from those under consideration) that are fixed-parameter tractable. Our trichotomy theorem is in fact derived by invoking two theorems:

- A new theorem showing that, for each set of existential positive queries, there exists a set of primitive positive queries such that the two sets exhibit the same complexity behavior (see Theorem 3.1). This new theorem, which we call the *equivalence theorem*, can be conceived of as the primary technical contribution of this article.
- A previously presented trichotomy on primitive positive queries [CM14a, CM15] (discussed in Section 2.4.)

1.3 Related work

The statement of our new trichotomy theorem generalizes, unifies, and strengthens a number of existing parameterized complexity classification results in the literature, namely:

- The dichotomy for model checking primitive positive formulas [Gro07], which built on a previous dichotomy [GSS01]; see also [CM14b].
- The dichotomy for model checking existential positive formulas [Che14a].
- The dichotomy for counting answers to quantifier-free primitive positive formulas [DJ04] (phrased as the problem of counting homomorphisms between relational structures).
- The trichotomy for counting answers to primitive positive formulas [CM14a, CM15], which trichotomy built on the previous work [DM13].

Let us emphasize that we only claim to generalize the parameterized complexity versions of the presented results. In some of the above works, such as the works [Gro07] and [DJ04], the problems that are classified as fixed-parameter tractable are also polynomial-time tractable. We can further remark that there are problems from the dichotomy theorem on model checking existential positive formulas [Che14a] that are shown to be fixed-parameter tractable but also NP-complete.

The techniques used to prove our equivalence theorem are algebraic and combinatorial, and are quite different in nature from and contrast with those used to prove the previous classifications, which were more graph-theoretic and logical in flavor. Indeed, while the graph-theoretic measure of *treewidth* played

[2] It can be phrased in terms of complexity classes: FPT $\neq$ W[1].

a key role in the statement and proof of the previous trichotomy as well as of the previous dichotomies on primitive positive queries, it is not at all needed to prove our equivalence theorem. To establish the equivalence theorem, we make a key application of the inclusion-exclusion counting principle to translate an existential positive formula to a finite set of primitive positive formulas (see Section 5.3), which, in the setup considered by the article, is crucial to handling and understanding disjunction. We believe that the developed theory that supports said application should provide a valuable foundation for coping with disjunction in logics that are more expressive than the one considered here.

2 Preliminaries

2.1 Basic definitions and notions

Note that $\cdot$ is sometimes used for multiplication of real numbers.

Polynomials. We remind the reader of some basic facts about polynomials which we will use throughout the paper. Here, a univariate polynomial p in a variable x is a function $p(x) = \sum_{i=0}^{d} a_i x^i$ where $d \geqslant 0$, each $a_i \in \mathbb{R}$ and $a_d \neq 0$, or the *zero polynomial* $p(x) = 0$. The a_i are called *coefficients* of p. The degree of a polynomial is defined as $-\infty$ in the case of the zero polynomial, and as d otherwise. Let $(x_0, y_0), \ldots, (x_n, y_n)$ be $n+1$ pairs of real numbers. Then there is a uniquely determined polynomial of degree at most n such that $p(x_i) = y_i$ for each i; consequently, a polynomial p of degree n that has at least $n+1$ zeroes (where a *zero* is a value x such that $p(x) = 0$) is the zero polynomial. If all x_i and y_i are rational numbers, then the coefficients a_i of this polynomial are rational numbers as well; moreover, the a_i can be computed in polynomial time.

Logic. We assume basic familiarity with the syntax and semantics of first-order logic. In this article, we focus on relational first-order logic where equality is not built-in to the logic. Hence, each *vocabulary/signature* under discussion consists only of relation symbols. We assume structures under discussion to be *finite* (that is, have finite universe); nonetheless, we sometimes describe structures as *finite* for emphasis. We assume that the relations of structures are represented as lists of tuples. We use the letters $\mathbf{A}$, $\mathbf{B}$, ... to denote structures, and the corresponding letters A, B, ... to denote their respective universes. When τ is a signature, we use $\mathbf{I}_\tau$ to denote the τ-structure with universe $\{a\}$ and where each relation symbol $R \in \tau$ has $R^{\mathbf{I}} = \{(a, \ldots, a)\}$. When $\mathbf{A}, \mathbf{B}$ are structures over the same signature τ, a *homomorphism* from $\mathbf{A}$ to $\mathbf{B}$ is a mapping $h : A \to B$ such that, for each $R \in \tau$ and each tuple $(a_1, \ldots, a_k) \in R^{\mathbf{A}}$, it holds that $(h(a_1), \ldots, h(a_k)) \in R^{\mathbf{B}}$.

We use the term *fo-formula* to refer to a first-order formula. An *ep-formula* (short for *existential positive formula*) is a fo-formula built from *atoms* (by which we refer to predicate applications of the form $R(v_1, \ldots, v_k)$, where R is a relation symbol and the v_i are variables), conjunction ($\wedge$), disjunction ($\vee$), and existential quantification ($\exists$). A *pp-formula* (short for *primitive positive formula*) is defined as an ep-formula where disjunction does not occur. An fo-

formula is *prenex* if it has the form $Q_1 v_1 \dots Q_n v_n \theta$ where θ is quantifier-free, that is, if all quantifiers occur in the front of the formula. The set of free variables of a formula ϕ is denoted by $\text{free}(\phi)$ and is defined as usual; a formula ϕ is a *sentence* if $\text{free}(\phi) = \varnothing$.

We now present some definitions and conventions that are not totally standard. A primary concern in this article is in counting satisfying assignments of fo-formulas on a finite structure. The count is sensitive to the set of variables over which assignments are considered; and, we will sometimes (but not always) want to count relative to a set of variables that is strictly larger than the set of free variables. Hence, we will often associate with each fo-formula ϕ a set V of variables called the *liberal variables*, denoted by $\text{lib}(\phi)$, which is required to be a superset of $\text{free}(\phi)$, that is, we require $\text{lib}(\phi) \supseteq \text{free}(\phi)$. Note that $\text{lib}(\phi)$ may contain variables that do not occur at all in atoms of ϕ. To indicate that V is the set of liberal variables of ϕ, we often use the notation $\phi(V)$; we also use $\phi(v_1, \dots, v_n)$, where the v_i are a listing of the elements of V. Relative to a formula $\phi(V)$, when $\mathbf{B}$ is a structure, we will use $\phi(\mathbf{B})$ to denote the set of assignments $f : V \to B$ such that $\mathbf{B}, f \models \phi$. *We assume that, in each prenex formula with liberal variables associated with it, no variable is both liberal and quantified.* We call an fo-formula ϕ *free* if $\text{free}(\phi) \neq \varnothing$, and *liberal* if $\text{lib}(\phi)$ is defined and $\text{lib}(\phi) \neq \varnothing$.

Example 2.1 Let us consider the formula $\phi(x, y, z) = R(x, y) \vee S(y, z)$. As indicated above, the notation $\phi(x, y, z)$ is used to indicate that $\text{lib}(\phi) = \{x, y, z\}$. As $\text{free}(\phi) = \{x, y, z\}$, we have $\text{lib}(\phi) = \text{free}(\phi)$. Define $\psi(x, y, z) = R(x, y)$ and $\psi'(x, y, z) = S(y, z)$. By the notation $\psi(x, y, z)$, we indicate that $\text{lib}(\psi) = \{x, y, z\}$; likewise, it holds that $\text{lib}(\psi') = \{x, y, z\}$. Notice that $\text{free}(\psi) = \{x, y\}$, so we have that $\text{lib}(\psi)$ is a proper superset of $\text{free}(\psi)$; in fact, the variable $z \in \text{lib}(\psi)$ does not occur at all in an atom of ψ. Define also $\theta(x, y) = R(x, y)$; by the notation $\theta(x, y)$, we indicate that $\text{lib}(\theta) = \{x, y\}$.

Observe that, for any structure $\mathbf{B}$, we have $\phi(\mathbf{B}) = \psi(\mathbf{B}) \cup \psi'(\mathbf{B})$ (and hence $|\phi(\mathbf{B})| = |\psi(\mathbf{B}) \cup \psi'(\mathbf{B})|$). Observe, however, that for any structure $\mathbf{B}$ where $\theta(\mathbf{B})$ is non-empty, it does not hold that $\phi(\mathbf{B}) = \theta(\mathbf{B}) \cup \psi'(\mathbf{B})$, since $\phi(\mathbf{B})$ contains only assignments defined on $\text{lib}(\phi) = \{x, y, z\}$, whereas $\theta(\mathbf{B})$ contains only assignments defined on $\text{lib}(\theta) = \{x, y\}$. □

pp-formulas. It is well-known [CM77] that there is a correspondence between prenex pp-formulas and relational structures. In particular, each prenex pp-formula $\phi(S)$ (on signature τ) with $\text{lib}(\phi) = S$ may be viewed as a pair $(\mathbf{A}, S)$ consisting of a structure $\mathbf{A}$ (on τ) and a set S; the universe A of $\mathbf{A}$ is the union of S with the variables appearing in ϕ, and the following condition defines the relations of $\mathbf{A}$: for each $R \in \tau$, a tuple $(a_1, \dots, a_k) \in A^k$ is in $R^{\mathbf{A}}$ if and only if $R(a_1, \dots, a_k)$ appears in ϕ. In the other direction, such a pair $(\mathbf{A}, S)$ can be viewed as a prenex pp-formula $\phi(S)$ where all variables in $A \backslash S$ are quantified and the atoms of ϕ are defined according to the above condition. A basic known fact [CM77] that we will use is that when $\phi(S)$ is a pp-formula corresponding to the pair $(\mathbf{A}, S)$, $\mathbf{B}$ is an arbitrary structure, and $f : S \to B$ is an arbitrary map,

it holds that $\mathbf{B}, f \models \phi(S)$ if and only if there is an extension f' of f that is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. *We will freely interchange between the structure view and the usual notion of a prenex pp-formula.* For a pp-formula specified as a pair $(\mathbf{A}, S)$, we typically assume that $S \subseteq A$.

Example 2.2 Consider the pp-formula $\phi(x, x', y, z) = \exists y' \exists u \exists v \exists w (E(x, x') \wedge E(y, y') \wedge F(u, v) \wedge G(u, w))$. The notation $\phi(x, x', y, z)$ indicates that $\mathrm{lib}(\phi) = \{x, x', y, z\}$. Note that $\mathrm{free}(\phi) = \{x, x', y\}$. To convert ϕ to a structure $\mathbf{A}$, we take the universe A of $\mathbf{A}$ to be the union of $\mathrm{lib}(\phi)$ with all variables appearing in ϕ, so $A = \{x, x', y, z, y', u, v, w\}$. Then, we define the relations as just described above, so $E^{\mathbf{A}} = \{(x, x'), (y, y')\}$, $F^{\mathbf{A}} = \{(u, v)\}$, and $G^{\mathbf{A}} = \{(u, w)\}$. The resulting pair representation of ϕ is $(\mathbf{A}, \{x, x', y, z\})$. □

Two structures are *homomorphically equivalent* if each has a homomorphism to the other. A structure is a *core* if it is not homomorphically equivalent to a proper substructure of itself. A structure $\mathbf{B}$ is a *core* of a structure $\mathbf{A}$ if $\mathbf{B}$ is a substructure of $\mathbf{A}$ that is a core and is homomorphically equivalent to $\mathbf{A}$. It is known that all cores of a structure are isomorphic and hence one sometimes speaks of *the* core of a structure.

For a prenex pp-formula $(\mathbf{A}, S)$ on signature τ, we define its *augmented structure*, denoted by $\mathsf{aug}(\mathbf{A}, S)$, to be the structure over the expanded vocabulary $\tau \cup \{R_a \mid a \in S\}$ (understood to be a disjoint union) where $R_a^{\mathsf{aug}(\mathbf{A},S)} = \{a\}$; we define the *core* of the pp-formula $(\mathbf{A}, S)$ to be the core of $\mathsf{aug}(\mathbf{A}, S)$.

The following fundamental facts on pp-formulas will be used throughout.

Theorem 2.3 (follows from [CM77]) *Suppose that each of the pairs $(\mathbf{A}, V)$, $(\mathbf{B}, V)$ is a prenex pp-formula. The formula $(\mathbf{B}, V)$ logically entails the formula $(\mathbf{A}, V)$ if and only if there exists a homomorphism from the structure $\mathsf{aug}(\mathbf{A}, V)$ to the structure $\mathsf{aug}(\mathbf{B}, V)$. The formulas $(\mathbf{A}, V)$, $(\mathbf{B}, V)$ are logically equivalent if and only if they have isomorphic cores, or equivalently, when $\mathsf{aug}(\mathbf{A}, V)$ and $\mathsf{aug}(\mathbf{B}, V)$ are homomorphically equivalent.*

ep-formulas. In order to discuss ep-formulas, we will employ the following terminology. An ep-formula is *disjunctive* if it is the disjunction of prenex pp-formulas; when ϕ is a disjunctive ep-formula with $\mathrm{lib}(\phi)$ defined, we typically assume that each of the pp-formulas ψ that appear as disjuncts of ϕ has $\mathrm{lib}(\psi)$ defined as $\mathrm{lib}(\phi)$. (In this way, for an arbitrary finite structure $\mathbf{B}$, it holds that $|\phi(\mathbf{B})| = |\bigcup_\psi \psi(\mathbf{B})|$, where the union is over all such disjuncts ψ.) An ep-formula is *all-free* if it is disjunctive and each pp-formula appearing as a disjunct is free. An ep-formula $\phi(S)$ is *normalized* if it is disjunctive and for each sentence disjunct $(\mathbf{A}, S)$ and any other disjunct $(\mathbf{A}', S)$, there is no homomorphism from $\mathsf{aug}(\mathbf{A}, S)$ to $\mathsf{aug}(\mathbf{A}', S)$ (equivalently, there is no homomorphism from $\mathbf{A}$ to $\mathbf{A}'$). It is straightforward to verify that there is an algorithm that, given an ep-formula, outputs a logically equivalent normalized ep-formula.

Graphs. To every prenex pp-formula $(\mathbf{A}, S)$ we assign a graph whose vertex set is $A \cup S$ and where two vertices are connected by an edge if they appear

together in a tuple of a relation of $\mathbf{A}$. A prenex pp-formula $(\mathbf{A}, S)$ is called *connected* if its graph is connected. A prenex pp-formula $(\mathbf{A}', S')$ is a *component* of a prenex pp-formula $(\mathbf{A}, S)$ over the same signature τ if there exists a set C that forms a connected component of the graph of $(\mathbf{A}, S)$, where:

- $S' = S \cap C$.
- For each relation $R \in \tau$, a tuple $(a_1, \ldots, a_k)$ is in $R^{\mathbf{A}'}$ if and only if $(a_1, \ldots, a_k) \in R^{\mathbf{A}} \cap C^k$.

Note that when this holds, the graph of $(\mathbf{A}', S')$ is the connected component of the graph of $(\mathbf{A}, S)$ on vertices C. We will use the fact that, if $\phi(V)$ is a prenex pp-formula and $\phi_1(V_1), \ldots, \phi_c(V_c)$ is a list of its components, then for any finite structure $\mathbf{B}$, it holds that $|\phi(\mathbf{B})| = \prod_{i=1}^{c} |\phi_i(\mathbf{B})|$.

Example 2.4 Let ϕ be the free prenex pp-formula from Example 2.2, and let $(\mathbf{A}, S)$ be the pair representation given there. The connected components of the graph of $(\mathbf{A}, S)$ are $\{x, x'\}$, $\{y, y'\}$, $\{z\}$, and $\{u, v, w\}$. There are thus four components of the formula $(\mathbf{A}, S)$; they are $(\mathbf{A}'_{\{x,x'\}}, \{x, x'\})$, $(\mathbf{A}'_{\{y,y'\}}, \{y\})$, $(\mathbf{A}'_{\{z\}}, \{z\})$, and $(\mathbf{A}'_{\{u,v,w\}}, \varnothing)$ (respectively), where each $\mathbf{A}'_C$ is the structure $\mathbf{A}'$ defined above, with respect to the set C.

Written logically, these four components are $\psi_1(x, x') = E(x, x')$, $\psi_2(y) = \exists y' E(y, y')$, $\psi_3(z) = \top$, and $\psi_4(\varnothing) = \exists u \exists v \exists w (F(u, v) \wedge G(u, w))$, respectively. Here, $\top$ denotes the empty conjunction (considered to be true). □

2.2 Counting complexity

Throughout, we use Σ to denote an alphabet over which strings are formed. All problems to be considered are viewed as counting problems. So, a *problem* is a mapping $Q : \Sigma^* \to \mathbb{N}$. We view decision problems as problems where, for each $x \in \Sigma^*$, it holds that $Q(x)$ is equal to 0 or 1. A *parameterization* is a mapping $\kappa : \Sigma^* \to \Sigma^*$. A parameterized problem is a pair (Q, κ) consisting of a problem Q and a parameterization κ. Throughout, by π_i we denote the operator that projects a tuple onto its ith coordinate.

A partial function $T : \Sigma^* \to \mathbb{N}$ is *polynomial-multiplied* with respect to a parameterization κ if there exists a computable function $f : \Sigma^* \to \mathbb{N}$ and a polynomial $p : \mathbb{N} \to \mathbb{N}$ such that, for each $x \in \mathrm{dom}(T)$, it holds that $T(x) \leqslant f(\kappa(x))p(|x|)$.

We now give a definition of FPT-computability for partial mappings.

Definition 2.5 *Let $\kappa : \Sigma^* \to \Sigma^*$ be a parameterization. A partial mapping $r : \Sigma^* \to \Sigma^*$ is* FPT-computable *with respect to κ if there exist a polynomial-multiplied function $T : \Sigma^* \to \mathbb{N}$ (with respect to κ) with $\mathrm{dom}(T) = \mathrm{dom}(r)$ and an algorithm A such that, for each string $x \in \mathrm{dom}(r)$, the algorithm A computes $r(x)$ within time $T(x)$; when this holds, we also say that r is* FPT-computable *with respect to κ* via *A.*

As is standard, we may and do freely interchange among elements of Σ^*, $\Sigma^* \times \Sigma^*$, and $\mathbb{N}$. We define FPT to be the class that contains a parameterized problem (Q, κ) if and only if Q is FPT-computable with respect to κ.

We now introduce a notion of reduction for counting problems, which is a form of Turing reduction. We use $\wp_{\mathsf{fin}}(A)$ to denote the set containing all finite subsets of A.

Definition 2.6 *A* counting FPT-reduction *from a parameterized problem (Q, κ) to another (Q', κ') consists of a computable function $h : \Sigma^* \to \wp_{\mathsf{fin}}(\Sigma^*)$, and an algorithm A such that:*

- *on an input x, A may make oracle queries of the form $Q'(y)$ with $\kappa'(y) \in h(\kappa(x))$, and*
- *Q is FPT-computable with respect to κ via A.*

We use Clique to denote the decision problem where (k, G) is a yes-instance when G is a graph that contains a clique of size $k \in \mathbb{N}$. By #Clique we denote the problem of counting, given (k, G), the number of k-cliques in the graph G. The parameterized versions of these problems, denoted by p-Clique and p-#Clique, are defined via the parameterization $\pi_1(k, G) = k$.

2.3 Counting case complexity

We employ the framework of *case complexity* to develop some of our complexity results. We present the needed elements of this framework for counting problems. The definitions and results here are due to [CM14a, CM15], are based on the theory of [Che14b], and are presented here for the sake of self-containment; see those articles for further discussion and motivation of the framework.

The case complexity framework was developed to prove results on restricted versions of parameterized problems where not all values of the parameter are permitted. This type of restricted problem arises naturally in query answering problems, where one often restricts the queries that are admissible, as is done here (for other examples, see [DJ04, Gro07, Che14b]).

The case complexity framework provides a notion of *case problem* and a notion of reduction between case problem. A case problem was originally [Che14b] defined as a language Q of pairs (that is, a subset of $\Sigma^* \times \Sigma^*$) where the first element of each pair is ultimately viewed as the parameter, along with a set $S \subseteq \Sigma^*$ restricting the permitted parameter values. In this article, as we are dealing with counting complexity, in lieu of considering languages, we will consider mappings $\Sigma^* \times \Sigma^* \to \mathbb{N}$. (Of course, a language of pairs can be naturally viewed as such a mapping by taking its characteristic function.)

One benefit of the framework is that the notion of reduction does not rely on any form of computability assumption on the sets S involved. Thus, in comparing case problems using this notion of reduction, one does not need to discuss the computability status of these sets S, even though in general, it is

usual that authors ultimately assume some form of computability on these sets (typically computable enumerability or computability).

Let us turn to the formal presentation of the framework. A *case problem* consists of a problem $Q : \Sigma^* \times \Sigma^* \to \mathbb{N}$ and a subset $S \subseteq \Sigma^*$, and is denoted $Q[S]$. Note that, although a problem above is defined as a mapping from Σ^* to $\mathbb{N}$, here we work with a problem that is a mapping from $\Sigma^* \times \Sigma^*$ to $\mathbb{N}$; this is natural in the current paper, where an input to the studied problem consists of two parts, a formula and a structure. Note that a mapping $\Sigma^* \times \Sigma^* \to \mathbb{N}$ can be naturally viewed as a mapping $\Sigma^* \to \mathbb{N}$, as there are natural and well-known ways to encode the elements of $\Sigma^* \times \Sigma^*$ as elements of Σ^*. For each case problem $Q[S]$, we define $\mathsf{param}\text{-}Q[S]$ as the parameterized problem (P, π_1) where $P(s, x)$ is defined as equal to $Q(s, x)$ if $s \in S$, and as 0 otherwise.

We have the following reduction notion for case problems.

Definition 2.7 *A* counting slice reduction *from a case problem $Q[S]$ to a second case problem $Q'[S']$ consists of*

- *a computably enumerable language $U \subseteq \Sigma^* \times \wp_{\mathsf{fin}}(\Sigma^*)$, and*
- *a partial function $r : \Sigma^* \times \wp_{\mathsf{fin}}(\Sigma^*) \times \Sigma^* \to \Sigma^*$ that has domain $U \times \Sigma^*$ and is FPT-computable with respect to (π_1, π_2) via an algorithm A that, on input (s, T, y), may make queries of the form $Q'(t, z)$ where $t \in T$,*

such that the following conditions hold:

- *(coverage) for each $s \in S$, there exists $T \subseteq S'$ such that $(s, T) \in U$, and*
- *(correctness) for each $(s, T) \in U$, it holds (for each $y \in \Sigma^*$) that $Q(s, y) = r(s, T, y)$.*

Let us provide some intuition for this definition. Here, when discussing an instance (s, y) of a case problem, we refer to the first part s as the parameter. The role of U is to provide all pairs (s, T) such that instances (of the first problem) with parameter s can be reduced to instances (of the second problem) whose parameters lie in T. Correspondingly, the *coverage* condition posits that each $s \in S$ is covered by the second set S' in the sense that there exists a pair $(s, T) \in U$ with $T \subseteq S'$. The partial function r is the actual reduction; given a pair $(s, T) \in U$ along with a string y, it computes the value $Q(s, y)$—this is what the *correctness* condition asserts. As here in this article we are dealing with counting complexity, we permit a form of Turing reduction; so, the algorithm A of the partial function r, upon being given a triple (s, T, y), may make (possibly multiple) queries to the second problem, so long as the queries are about instances whose parameter falls into T.

We have the following key property of counting slice reducibility.

Theorem 2.8 *[CM14a] Counting slice reducibility is transitive.*

The following theorem shows that, from a counting slice reduction, one can obtain complexity results for the corresponding parameterized problems.

Theorem 2.9 *[CM14a] Let $Q[S]$ and $Q'[S']$ be case problems. Suppose that $Q[S]$ counting slice reduces to $Q'[S']$, and that both S and S' are computable. Then* param-$Q[S]$ *counting FPT-reduces to* param-$Q'[S']$.

2.4 Classification of pp-formulas

We present the complexity classification of pp-formulas previously presented in [CM14a, CM15]. The following definitions are adapted from that article. Let $(\mathbf{A}, S)$ be a prenex pp-formula, let $\mathbf{D}$ be the core thereof, and let $G = (D, E)$ be the graph of $\mathbf{D}$. An $\exists$-*component* of $(\mathbf{A}, S)$ is a graph of the form $G[V']$ where there exists $V \subseteq D$ that is the vertex set of a component of $G[D \backslash S]$ and V' is the union of V with all vertices in S having an edge to V. Define $\mathsf{contract}(\mathbf{A}, S)$ to be the graph on vertex set S obtained by starting from $G[S]$ and adding an edge between any two vertices that appear together in an $\exists$-component of $(\mathbf{A}, S)$.

Let Φ be a set of prenex pp-formulas. Let us say that Φ satisfies the *contraction condition* if the graphs in the set $\mathsf{contract}(\Phi) := \{\mathsf{contract}(\phi) \mid \phi \in \Phi\}$ are of bounded treewidth. Let us say that Φ satisfies the *tractability condition* if it satisfies the contraction condition and, in addition, the cores of Φ are of bounded treewidth; here, the treewidth of a prenex pp-formula is defined as that of its graph. We omit the definition of treewidth, as it is both well-known and not needed to understand the main technical proof of this article (which is in Section 5).

Definition 2.10 *We define* count *to be the problem that maps a pair $(\phi(V), \mathbf{B})$ consisting of a fo-formula and a finite structure to the value $|\phi(\mathbf{B})|$.*

Theorem 2.11 *[CM14a] Let Φ be a set of prenex pp-formulas that satisfies the tractability condition. Then, the restriction of* param-$\mathsf{count}[\Phi]$ *to $\Phi \times \Sigma^*$ is an FPT-computable partial function.*

Theorem 2.12 *[CM14a] Let Φ be a set of prenex pp-formulas of bounded arity that does not satisfy the tractability condition.*

1. *If Φ satisfies the contraction condition, then it holds that* $\mathsf{count}[\Phi]$ *and* $\textsc{Clique}[\mathbb{N}]$ *are interreducible, under counting slice reductions.*
2. *Otherwise, there exists a counting slice reduction from $\#\textsc{Clique}[\mathbb{N}]$ to* $\mathsf{count}[\Phi]$.

We say that a set of formulas Φ has *bounded arity* if there exists a constant $k \geqslant 1$ that upper bounds the arity of each relation symbol appearing in a formula in Φ.

3 Main theorems

The following theorem, which we call the *equivalence theorem* and which is proved in Section 5, is our primary technical result; it is used to derive our

complexity trichotomy on ep-formulas from the known complexity trichotomy on pp-formulas (which was presented in Section 2.4).

Theorem 3.1 *(Equivalence theorem) Let Φ be a set of ep-formulas. There exists a set Φ^+ of prenex pp-formulas with the following property: the two counting case problems $\mathsf{count}[\Phi]$ and $\mathsf{count}[\Phi^+]$ are interreducible under counting slice reductions. In particular, there exists an algorithm that computes, given an ep-formula ϕ, a finite set ϕ^+ of prenex pp-formulas such that for any set Φ of ep-formulas, the set Φ^+ defined as $\bigcup\{\phi^+ \mid \phi \in \Phi\}$ has the presented property.*

We now state our trichotomy theorem on the complexity of counting answers to ep-formulas, and show how to prove it using the equivalence theorem.

Theorem 3.2 *(Trichotomy theorem) Let Φ be a computable set of ep-formulas of bounded arity, and let Φ^+ be the set of pp-formulas given by Theorem 3.1.*

1. *If Φ^+ satisfies the tractability condition, then it holds that $\mathsf{param\text{-}count}[\Phi]$ is in FPT.*
2. *If Φ^+ does not satisfy the tractability condition but satisfies the contraction condition, then it holds that $\mathsf{param\text{-}count}[\Phi]$ is interreducible with p-Clique under counting FPT-reduction.*
3. *Otherwise, there is a counting FPT-reduction from the problem p-#Clique to $\mathsf{param\text{-}count}[\Phi]$.*

Proof. For (1), we use the counting slice reduction (U, r) from $\mathsf{count}[\Phi]$ to $\mathsf{count}[\Phi^+]$ given by Theorem 3.1. In particular, given as input $(\phi, \mathbf{B})$, it is first checked if $\phi \in \Phi$; if not, 0 is output. Otherwise, the algorithm for r is invoked on $(\phi, \phi^+, \mathbf{B})$, where ϕ^+ is as defined in the statement of Theorem 3.1; queries to $\mathsf{count}(\psi, \mathbf{B})$ where $\psi \in \Phi^+$ are resolved according to the algorithm of Theorem 2.11.

For (2) and (3), we make use of the result (Theorem 3.1) that the problems $\mathsf{count}[\Phi]$ and $\mathsf{count}[\Phi^+]$ are interreducible under counting slice reductions. For (2), we have from Theorem 2.12 that $\mathsf{count}[\Phi^+]$ and Clique$[\mathbb{N}]$ are interreducible under counting slice reductions. Hence, we obtain that the problems Clique$[\mathbb{N}]$ and $\mathsf{count}[\Phi]$ are interreducible under counting slice reductions, and the result follows from Theorem 2.9. For (3), we have from Theorem 2.12 that there is a counting slice reduction from #Clique$[\mathbb{N}]$ to $\mathsf{count}[\Phi^+]$, and hence from #Clique$[\mathbb{N}]$ to $\mathsf{count}[\Phi]$; the result then follows from Theorem 2.9. □

Let us remark that when case (2) applies, a consequence of this theorem is that the problem $\mathsf{param\text{-}count}[\Phi]$ is not in FPT unless W[1] is in FPT, since p-Clique is W[1]-complete; in a similar fashion, when case (3) applies, the problem $\mathsf{param\text{-}count}[\Phi]$ is not in FPT unless ♯W[1] is in FPT, since p-#Clique is ♯W[1]-complete.

4 Examples

Before proving the equivalence theorem in full generality, we discuss some example ep-formulas to illustrate and preview some of the issues and difficulties with which the argument needs to cope.

Example 4.1 Consider the formula

$$\phi(w,x,y,z) := E(x,y) \wedge (E(w,x) \vee (E(y,z) \wedge E(z,z))).$$

As a first simplification step, we bring disjunction to the outermost level in ϕ:

$$\begin{aligned}&\phi(w,x,y,z)\\ &\equiv (E(x,y) \wedge E(w,x)) \vee (E(x,y) \wedge E(y,z) \wedge E(z,z)).\end{aligned}$$

Now let us set $\phi_1(w,x,y,z) \equiv E(x,y) \wedge E(w,x)$ and also set $\phi_2(w,x,y,z) \equiv E(x,y) \wedge E(y,z) \wedge E(z,z)$. We can use inclusion-exclusion to count the number of satisfying assignments of ϕ on a structure $\mathbf{B}$ by

$$|\phi(\mathbf{B})| = |\phi_1(\mathbf{B})| + |\phi_2(\mathbf{B})| - |(\phi_1 \wedge \phi_2)(\mathbf{B})|.$$

One point to observe is that, in this last expression, the count $|\phi_1(\mathbf{B})|$ needs to be determined with respect to its set of liberal variabes $\mathrm{lib}(\phi_1) = \{w,x,y,z\}$, even though z does not appear in any atom of ϕ_1. If the count $|\phi_1(\mathbf{B})|$ is not computed in this way, the above expression for $|\phi(\mathbf{B})|$ fails to hold in general. The situation is analogous for the formula ϕ_2, where w does not appear in any atom. □

Example 4.2 In general, if we are given an ep-formula $\phi = \phi_1 \vee \ldots \vee \phi_n$ where the ϕ_i are pp-formulas, then to compute the count $|\phi(\mathbf{B})|$ of ϕ relative to $\mathbf{B}$, it suffices to know the count for each of the $2^n - 1$ pp-formulas obtained by taking a conjunction of a non-empty subset of the ϕ_i. In this example, we will see that, in fact, one does not always need to consider all of these conjunctions. To this end, set $V = \{w,x,y,z\}$ and set

$$\phi(V) = \phi_1(V) \vee \phi_2(V) \vee \phi_3(V)$$

where $\phi_1(V) = E(x,y) \wedge E(y,z)$, $\phi_2(V) = E(z,w) \wedge E(w,x)$ and $\phi_3(V) = E(w,x) \wedge E(x,y)$. Applying inclusion-exclusion, we obtain

$$\begin{aligned}|\phi(\mathbf{B})| =& |\phi_1(\mathbf{B})| + |\phi_2(\mathbf{B})| + |\phi_3(\mathbf{B})|\\ &- |(\phi_1 \wedge \phi_2)(\mathbf{B})| - |(\phi_1 \wedge \phi_3)(\mathbf{B})|\\ &- |(\phi_2 \wedge \phi_3)(\mathbf{B})| + |(\phi_1 \wedge \phi_2 \wedge \phi_3)(\mathbf{B})|.\end{aligned}$$

Now observe that the formulas ϕ_1, ϕ_2 and ϕ_3 are actually equivalent to each other up to renaming variables; consequently, these formulas are equivalent in that, for any structure $\mathbf{B}$, they yield the same count: $|\phi_1(\mathbf{B})| = |\phi_2(\mathbf{B})| = |\phi_3(\mathbf{B})|$. In Section 5.1, we formalize and give a characterization of this notion

of equivalence (on pp-formulas). The formulas $\phi_1 \wedge \phi_3$ and $\phi_2 \wedge \phi_3$ are also equivalent in this sense. We may thus obtain the following expression for $|\phi(\mathbf{B})|$.

$$\begin{aligned}|\phi(\mathbf{B})| =& 3 \cdot |\phi_1(\mathbf{B})| - |(\phi_1 \wedge \phi_2)(\mathbf{B})| \\ & - 2 \cdot |(\phi_1 \wedge \phi_3)(\mathbf{B})| + |(\phi_1 \wedge \phi_2 \wedge \phi_3)(\mathbf{B})|.\end{aligned}$$

So far, we have only unified formulas that are equivalent up to renaming variables. In our parameterized complexity setting where ϕ is the parameter, this does not yield a significant decrease in the complexity of computing $|\phi(\mathbf{B})|$. However, we will now observe a simplification that is more substantial in this sense. Namely, one can verify that the formulas $\phi_1 \wedge \phi_2$ and $\phi_1 \wedge \phi_2 \wedge \phi_3$ are identical. So, if we identify their terms in this last expression for $|\phi(\mathbf{B})|$, we obtain a cancellation and arrive to the following expression:

$$|\phi(\mathbf{B})| = 3 \cdot |\phi_1(\mathbf{B})| - 2 \cdot |(\phi_1 \wedge \phi_3)(\mathbf{B})|.$$

The savings obtained by observing this cancellation are significant, in the following sense. The pp-formulas $\phi_1 \wedge \phi_2$ and $\phi_1 \wedge \phi_2 \wedge \phi_3$, which were cancelled, were the only formulas in the expression for $|\phi(\mathbf{B})|$ which did not have treewidth 1; they had treewidth 2. As it is known that the runtime of evaluation algorithms for quantifier-free pp-formulas scales with their treewidth [Mar10a], this reduction in treewidth yields a superior runtime for evaluating $|\phi(\mathbf{B})|$. □

As we have seen in the above examples, counting on an ep-formula can, via inclusion-exclusion, reduce to counting on a finite set of pp-formulas. (This is carried out in our argument; see Section 5.3). As just seen in Example 4.2, there can be some subtlety in choosing a desirable set of pp-formulas to reduce to. One question not addressed so far is how one can reduce from counting on a such obtained set of pp-formulas to counting on the original ep-formula. To this end, let us revisit our first example.

Example 4.3 Let us consider again the formulas of Example 4.1. Assume that we are given access to an oracle that lets us compute $|\phi(\mathbf{D})|$, for any structure $\mathbf{D}$ of our choice. We will see that, given a structure $\mathbf{B}$, we can compute $|\phi_1(\mathbf{B})|$, $|\phi_2(\mathbf{B})|$, and $|(\phi_1 \wedge \phi_2)(\mathbf{B})|$ efficiently using this oracle.

To see this, consider the structure $\mathbf{C}$ with universe $C = \{1, 2, 3, 4\}$ and $E^{\mathbf{C}} = \{(1, 2), (2, 3), (3, 4), (4, 4)\}$. It is easy to check that the formulas ϕ_1, ϕ_2 and $\phi_1 \wedge \phi_2$ all have a different number of answers with respect to $\mathbf{C}$. Now note that for every pp-formula ψ and every pair of structures $\mathbf{D}_1$, $\mathbf{D}_2$ we have $|\psi(\mathbf{D}_1 \times \mathbf{D}_2)| = |\psi(\mathbf{D}_1)| \cdot |\psi(\mathbf{D}_2)|$. Querying the oracle for $|\phi(\cdot)|$ on $\mathbf{B} \times \mathbf{C}^i$ for the values $i = 0, 1, 2$, we obtain the linear system

$$A \begin{pmatrix} |\phi_1(\mathbf{B})| \\ |\phi_2(\mathbf{B})| \\ -|(\phi_1 \wedge \phi_2)(\mathbf{B})| \end{pmatrix} = \begin{pmatrix} (\phi(\mathbf{B}) \\ \phi(\mathbf{B} \times \mathbf{C}) \\ \phi(\mathbf{B} \times \mathbf{C}^2) \end{pmatrix}$$

with

$$A = \begin{pmatrix} 1 & 1 & 1 \\ |\phi_1(\mathbf{C})| & |\phi_2(\mathbf{C})| & |(\phi_1 \wedge \phi_2)(\mathbf{C})| \\ |\phi_1(\mathbf{C})|^2 & |\phi_2(\mathbf{C})|^2 & |(\phi_1 \wedge \phi_2)(\mathbf{C})|^2 \end{pmatrix}.$$

Note that the entries of A can be computed efficiently, and the vector on the right-hand-side of the equation can be provided by our oracle. The matrix A is a Vandermonde matrix, as a consequence of the choice of $\mathbf{C}$. Thus, the system has a unique solution and can be solved to determine $|\phi_1(\mathbf{B})|$, $|\phi_2(\mathbf{B})|$, and $|(\phi_1 \wedge \phi_2)(\mathbf{B})|$, as desired. □

In Example 4.3 we have seen that, for the particular ep-formula ϕ discussed, counting on ϕ is in a certain sense interreducible with counting on the pp-formulas

$$\{\phi_1, \phi_2, \phi_1 \wedge \phi_2\}.$$

The statement of the equivalence theorem (Theorem 3.1) asserts that for *any* ep-formula ϕ, there exists a finite set ϕ^+ of pp-formulas such that one has this interreducibility.

5 Proof of equivalence theorem

In this section, we give a decidable characterization of *counting equivalence* (Section 5.1); we then study a relaxation thereof which we call *semi-counting equivalence* (Section 5.2); we prove the equivalence theorem in the particular case of all-free ep-formulas (Section 5.3); and, we end by proving the equivalence theorem in its full generality (Section 5.4). Throughout this section, we generally assume pp-formulas to be prenex.

5.1 Counting equivalence

As we have seen in the examples of Section 4, it will be important to see when two different pp-formulas give same number of answers for every structure, because it will allow us to make simplifications in formulas we get by inclusion-exclusion. To this end, we make the following definition.

Definition 5.1 *Define two fo-formulas $\phi(V), \phi'(V')$ to be* counting equivalent *if they are over the same vocabulary τ and for each finite τ-structure $\mathbf{B}$ it holds that $|\phi(\mathbf{B})| = |\phi'(\mathbf{B})|$.*

In this subsection, we characterize counting equivalence for pp-formulas. To approach the characterization, we start off with an example.

Example 5.2 It is apparent that logically equivalent formulas are counting equivalent, but the converse direction is not true. To see this, consider the pp-formulas $\phi_1(x, y) = E(x, y)$ and $\phi_2(w, z) = E(w, z)$. Obviously, ϕ_1 and ϕ_2 are counting equivalent (they just count the number of tuples in the relation E of a structure $\mathbf{B}$). But ϕ_1 and ϕ_2 are not logically equivalent; indeed, the assignments in $\phi_1(\mathbf{B})$ and $\phi_2(\mathbf{B})$ assign values to different variables.

Note that one way of witnessing the counting equivalence of ϕ_1 and ϕ_2 is simply renaming the variable w to x and z to y to get equivalent formulas. Since this syntactic renaming obviously does not change the number of satisfying assignments, one can conclude that ϕ_1 and ϕ_2 are counting equivalent. □

Example 5.2 motivates the following definition.

Definition 5.3 *We say that two pp-formulas*

$$(\mathbf{A}, S), (\mathbf{A}', S')$$

over the same signature are renaming equivalent *if there exist surjections* $h : S \to S'$ *and* $h' : S' \to S$ *that can be extended to homomorphisms* $\bar{h} : \mathbf{A} \to \mathbf{A}'$ *and* $\bar{h}' : \mathbf{A}' \to \mathbf{A}$*, respectively.*

Informally speaking, on pp-formulas, two formulas are renaming equivalent if they become logically equivalent after a renaming of variables, as occurred in Example 5.2. Hence, renaming equivalence is a relaxation of logical equivalence. Recall that logical equivalence of pp-formulas was characterized, in Theorem 2.3.

The main theorem of this subsection is that renaming equivalence does not only imply counting equivalence but is actually equivalent to it.

Theorem 5.4 *Two pp-formulas*

$$\phi_1(S_1), \phi_2(S_2)$$

are counting equivalent if and only if they are renaming equivalent.

Note that Theorem 5.4 gives a syntactic/algebraic characterization of counting equivalence which makes counting equivalence decidable by a straightforward algorithm and in fact even puts it into NP.

Before we prove Theorem 5.4, we start off with an simple observation that will be helpful in the proof.

Observation 5.5 *Let* ϕ *and* ϕ' *be counting equivalent pp-formulas. Then* $|\mathrm{lib}(\phi)| = |\mathrm{lib}(\phi')|$.

Proof. Let $\mathbf{C}$ be a structure that interprets every relation symbol in R of ϕ by $R^{\mathbf{C}} := \{0,1\}^{\mathsf{arity}(R)}$. Then $|\phi(\mathbf{C})| = 2^{|\mathrm{lib}(\phi)|}$ and $|\phi'(\mathbf{C})| = 2^{|\mathrm{lib}(\phi')|}$ and the claim follows directly. □

Proof. (Theorem 5.4) We begin with the backward direction; let $h_1 : S_1 \to S_2$ and $h_2 : S_2 \to S_1$ be the surjections from the definition of renaming equivalence. The existence of these surjections implies that $|S_1| = |S_2|$ and that each of h_1, h_2 is a bijection. Let $\mathbf{B}$ be an arbitrary structure. For each $f : S_2 \to B$ in $\phi_2(\mathbf{B})$, it is straightforward to verify that the composition $f(h_1)$ is in $\phi_1(\mathbf{B})$. Since the mapping that takes each such f to $f(h_1)$ is injective (due to h_1 being a bijection), we obtain that $|\phi_1(\mathbf{B})| \geqslant |\phi_2(\mathbf{B})|$. By symmetric reasoning, we can obtain that $|\phi_1(\mathbf{B})| \leqslant |\phi_2(\mathbf{B})|$, and we conclude that $|\phi_1(\mathbf{B})| = |\phi_2(\mathbf{B})|$.

For the other direction, let $\phi_1(S_1)$ and $\phi_2(S_2)$ be two pp-formulas over a common vocabulary τ that are not renaming equivalent; let $(\mathbf{A}_1, S_1)$ and $(\mathbf{A}_2, S_2)$ be the corresponding structures. By way of contradiction, assume that ϕ_1 and ϕ_2 are counting equivalent. If it holds that $|\mathrm{lib}(\phi_1)| \neq |\mathrm{lib}(\phi_2)|$, we are done by

Observation 5.5. So we may assume, after potentially renaming some variables, that $\mathrm{lib}(\phi_1) = \mathrm{lib}(\phi_2) =: S$.

When $\mathbf{C}$, $\mathbf{D}$ are structures with $S \subseteq C \cap D$, let us define $\hom(\mathbf{C}, \mathbf{D}, S)$ to be the set of mappings from S to D that can be extended to a homomorphism from $\mathbf{C}$ to $\mathbf{D}$; denote by $\mathrm{surj}(\mathbf{C}, \mathbf{D}, S)$ the surjections $h : S \to S$ that lie in $\hom(\mathbf{C}, \mathbf{D}, S)$.

As $(\mathbf{A}_1, S_1)$ and $(\mathbf{A}_2, S_2)$ are by hypothesis not renaming equivalent, we may assume, without loss of generality, that $\mathrm{surj}(\mathbf{A}_1, \mathbf{A}_2, S) = \varnothing$. For $T \subseteq S$ let us use $\hom_T(\mathbf{A}_1, \mathbf{A}_2, S)$ to denote the set of mappings $h \in \hom(\mathbf{A}_1, \mathbf{A}_2, S)$ such that $h(S) \subseteq T$. By inclusion-exclusion we get

$$|\mathrm{surj}(\phi_1, \phi_2, S)| = \sum_{T \subseteq S} (-1)^{|S|-|T|} |\hom_T(\mathbf{A}_1, \mathbf{A}_2, S)|.$$

For $i \geqslant 0$ let $\hom_{i,T}(\mathbf{A}_1, \mathbf{A}_2, S)$ be the set of mappings $h \in \hom(\mathbf{A}_1, \mathbf{A}_2, S)$ such that h maps exactly i variables from S into T. Now for each $j = 1, \ldots, |S|$ we construct a new structure $\mathbf{D}_{j,T}$ over the domain $D_{j,T}$. To this end, let $a^{(1)}, \ldots, a^{(j)}$ be copies of $a \in T$ that are not in A_2. Then we set

$$D_{j,T} := \{a^{(k)} \mid a \in A_2, a \in T, k \in [j]\} \cup (A_2 \backslash T).$$

We define a mapping $B : A_2 \to \mathcal{P}(D_{j,T})$, where $\mathcal{P}(D_{j,T})$ is the power set of $D_{j,T}$, by

$$B(a) := \begin{cases} \{a^{(k)} \mid k \in [j]\}\}, & \text{if } a \in T \\ \{a\}, & \text{otherwise.} \end{cases}$$

For every relation symbol $R \in \tau$ we define

$$R^{\mathbf{D}_{T,j}} := \bigcup_{(d_1, \ldots, d_s) \in R^{\mathbf{A}_2}} B(d_1) \times \ldots \times B(d_s).$$

Then every $h \in \hom_{i,T}(\mathbf{A}_1, \mathbf{A}_2, S)$ corresponds to j^i mappings in $\hom(\mathbf{A}_1, \mathbf{D}_{j,T}, S)$. Thus for each j we get

$$\sum_{i=1}^{|S|} j^i |\hom_{i,T}(\mathbf{A}_1, \mathbf{A}_2, S)| = |\hom(\mathbf{A}_1, \mathbf{D}_{j,T}, S)|.$$

This is a linear system of equations and the corresponding matrix is a Vandermonde matrix; consequently, the value $\hom_T(\mathbf{A}_1, \mathbf{A}_2, S) = \hom_{|S|,T}(\mathbf{A}_1, \mathbf{A}_2, S)$ can efficiently be computed from $|\hom(\mathbf{A}_1, \mathbf{D}, S)| = |\phi_1(\mathbf{D})|$ for some structures $\mathbf{D}$. We can similarly determine $|\hom_T(\mathbf{A}_2, \mathbf{D}, S)|$ as a function of $|\phi_2(\mathbf{D})|$ for the same structures $\mathbf{D}$. Since $|\phi_1(\mathbf{D})| = |\phi_2(\mathbf{D})|$ for every structure $\mathbf{D}$ by assumption, it follows that for every subset $T \subseteq S$ we have

$$|\hom_T(\mathbf{A}_1, \mathbf{A}_2, S)| = |\hom_T(\mathbf{A}_2, \mathbf{A}_2, S)|.$$

But then we have

$$|\mathrm{surj}(\mathbf{A}_1, \mathbf{A}_2, S)| = |\mathrm{surj}(\mathbf{A}_2, \mathbf{A}_2, S)|.$$

Since $\mathrm{surj}(\mathbf{A}_1, \mathbf{A}_2, S) = \varnothing$ and $\mathrm{id} \in \mathrm{surj}(\mathbf{A}_2, \mathbf{A}_2, S)$, this is a contradiction. Consequently, we obtain that ϕ_1 and ϕ_2 are not counting equivalent. □

5.2 Semi-counting equivalence

In this subsection, we study a relaxation of the notion of *counting equivalence*. This notion will be necessary when we emulate the approach of Example 4.3 in the proof of the Equivalence theorem: we will again construct a system of linear equations that we want to solve. In order to ensure solvability, we will make sure that the matrix of the system is again a Vandermonde matrix which in particular means that all its entries must be positive. Consequently, since the entries are of the form $|\phi(\mathbf{C})|^k$ for pp-formulas ϕ some carefully chosen structure $\mathbf{C}$ and integers k, it will be necessary to understand counting equivalence in the case where $\phi(\mathbf{C})$ is non-empty. The necessary notion is formalized by the following definition.

Definition 5.6 *Call two prenex pp-formulas*

$$\phi_1(V_1), \phi_2(V_2)$$

on the same vocabulary semi-counting equivalent *if for each finite structure* $\mathbf{B}$ *such that* $|\phi_1(\mathbf{B})| > 0$ *and* $|\phi_2(\mathbf{B})| > 0$*, it holds that* $|\phi_1(\mathbf{B})| = |\phi_2(\mathbf{B})|$.

Example 5.7 The pp-formulas $\phi_1(x, y) = E(x, y)$ and $\phi_2(x, y) = \exists z(E(x, y) \wedge F(z))$ are not counting equivalent, because for every structure $\mathbf{B}$ for which $F^{\mathbf{B}} = \varnothing$, we have $|\phi_2(\mathbf{B})| = 0$ while $|\phi_1(\mathbf{B})|$ may be non-zero if $E^{\mathbf{B}}$ is non-empty. But if we have for a structure $\mathbf{B}$ such that $|\phi_2(\mathbf{B})| > 0$, then $F^{\mathbf{B}} \neq \varnothing$ and it is straightforward to verify that $|\phi_1(\mathbf{B})| = |\phi_2(\mathbf{B})|$. Consequently, we have that ϕ_1 and ϕ_2 are semi-counting equivalent. □

For each free prenex pp-formula $\phi(V)$, define $\widehat{\phi}(V)$ to be the pp-formula obtained from ϕ by removing each atom that occurs in a non-liberal component of ϕ (a component of ϕ not having liberal variables).

Example 5.8 Consider the pp-formula ϕ discussed in Examples 2.2 and 2.4. This pp-formula has 4 components, namely, the pp-formulas $\psi_1(x, x')$, $\psi_2(y)$, $\psi_3(z)$, and $\psi_4(\varnothing)$ defined in Example 2.4. The formulas ψ_1, ψ_2, and ψ_3 are liberal, but the formula ψ_4 is not liberal. Recall that the formula $\phi(x, x', y, z)$ is equal to

$$\exists y' \exists u \exists v \exists w(E(x, x') \wedge E(y, y') \wedge F(u, v) \wedge G(u, w))$$

and that we have $\psi_4(\varnothing) = \exists u \exists v \exists w(F(u, v) \wedge G(u, w))$. We hence have that $\widehat{\phi}(x, x', y, z)$ is the formula

$$\exists y' \exists u \exists v \exists w(E(x, x') \wedge E(y, y')).$$

□

The following characterization of semi-counting equivalence is the main theorem of this subsection.

Theorem 5.9 *Let $\phi_1(V_1), \phi_2(V_2)$ be two free prenex pp-formulas. It holds that $\phi_1(V_1)$ and $\phi_2(V_2)$ are semi-counting equivalent if and only if $\widehat{\phi_1}(V_1)$ and $\widehat{\phi_2}(V_2)$ are counting equivalent.*

We will use the following proposition in the proof of Theorem 5.9.

Proposition 5.10 *Let $\phi(V)$ be a free prenex pp-formula. Then for every structure $\mathbf{B}$ we have $\phi(\mathbf{B}) = \varnothing$ or $\phi(\mathbf{B}) = \widehat{\phi}(\mathbf{B})$.*

Proof. Let $\mathbf{B}$ be a structure. Let ψ be the conjunction of the components deleted from ϕ to obtain $\widehat{\phi}$. If ψ is false on $\mathbf{B}$, then obviously $\phi(\mathbf{B}) = \varnothing$. Otherwise, ψ is true on $\mathbf{B}$, and for any assignment $f : V \to B$, it holds that $\mathbf{B}, f \models \phi$ if and only if $\mathbf{B}, f \models \widehat{\phi}$. □

Proof. (Theorem 5.9) Assume first that $\widehat{\phi_1}$ and $\widehat{\phi_2}$ are counting equivalent. Let $\mathbf{B}$ be a structure. Then if $|\phi_1(\mathbf{B})| > 0$ and $|\phi_2(\mathbf{B})| > 0$, we have by Proposition 5.10 and counting equivalence of $\widehat{\phi_1}$ and $\widehat{\phi_2}$ that $|\phi_1(\mathbf{B})| = |\widehat{\phi_1}(\mathbf{B})| = |\widehat{\phi_2}(\mathbf{B})| = |\phi_2(\mathbf{B})|$, so ϕ_1 and ϕ_2 are semi-counting equivalent.

For the other direction let now ϕ_1 and ϕ_2 be semi-counting equivalent. By way of contradiction, we assume that $\widehat{\phi_1}$ and $\widehat{\phi_2}$ are not counting equivalent. Then by definition there is a structure $\mathbf{B}$ such that $|\widehat{\phi_1}(\mathbf{B})| \neq |\widehat{\phi_2}(\mathbf{B})|$. Note that each component of $\widehat{\phi_1}$ and $\widehat{\phi_2}$ has a liberal variable.

Let $\mathbf{I} = \mathbf{I}_\tau$, where τ is the vocabulary of ϕ_1 and ϕ_2. For each $k \in \mathbb{N}$ we denote by $\mathbf{B} + k\mathbf{I}$ the structure we get from $\mathbf{B}$ by disjoint union with k copies of $\mathbf{I}$. Note that for $k > 0$ we have $|\phi(\mathbf{B} + k\mathbf{I})| > 0$ for every pp-formula ϕ. Consequently, for every $k > 0$ we have $|\phi_1(\mathbf{B} + k\mathbf{I})| = |\widehat{\phi_1}(\mathbf{B} + k\mathbf{I})|$ and $|\phi_2(\mathbf{B} + k\mathbf{I})| = |\widehat{\phi_2}(\mathbf{B} + k\mathbf{I})|$ by Proposition 5.10. By the semi-counting equivalence of ϕ_1 and ϕ_2 we also have $|\phi_1(\mathbf{B} + k\mathbf{I})| = |\phi_2(\mathbf{B} + k\mathbf{I})|$ for all $k > 0$. It follows that $|\widehat{\phi_1}(\mathbf{B} + k\mathbf{I})| = |\widehat{\phi_2}(\mathbf{B} + k\mathbf{I})|$ for $k > 0$.

Let $\phi_{1,1}, \ldots, \phi_{1,n}$ denote the components of $\widehat{\phi_1}$, and let $\phi_{2,1}, \ldots, \phi_{2,m}$ denote the components of $\widehat{\phi_2}$. Because every component of $\widehat{\phi_1}$ has a liberal variable, we have

$$\begin{aligned} |\widehat{\phi_1}(\mathbf{B} + k\mathbf{I})| &= \sum_{J \subseteq [n]} k^{n-|J|} \prod_{j \in J} |\phi_{1,j}(\mathbf{B})| \\ &= \sum_{\ell=0}^{n} k^{n-\ell} \sum_{J \subseteq [n], |J|=\ell} \prod_{j \in J} |\phi_{1,j}(\mathbf{B})|. \end{aligned}$$

We can express $|\widehat{\phi_2}(\mathbf{B} + k\mathbf{I})|$ analogously. The expressions are polynomials in k and they are equal for every positive integer k by the observations above; thus the coefficients of the polynomials must coincide. The coefficients of k^0, namely the values $\prod_{j \in [n]} |\phi_{1,j}(\mathbf{B})|$ and $\prod_{j \in [m]} |\phi_{2,j}(\mathbf{B})|$, are thus equal. But then we get

$$|\widehat{\phi_1}(\mathbf{B})| = \prod_{j \in [n]} |\phi_{1,j}(\mathbf{B})| = \prod_{j \in [m]} |\phi_{2,j}(\mathbf{B})| = |\widehat{\phi_2}(\mathbf{B})|,$$

which is a contradiction to our assumption. □

Corollary 5.11 *Semi-counting equivalence is an equivalence relation (on pp-formulas).*

We now present a lemma that will be of utility; it is proved by induction.

Lemma 5.12 *Let $\phi_1(S_1), \ldots, \phi_n(S_n)$ be pp-formulas over the same vocabulary τ, which are liberal (that is, with each $|S_i| > 0$). Then there is a structure $\mathbf{C}$ (over τ) such that*

- *for all pp-formulas ϕ (over τ) it holds that $|\phi(\mathbf{C})| > 0$, and*
- *for all $i, j \in [n]$ such that ϕ_i and ϕ_j are not semi-counting equivalent, it holds that $|\phi_i(\mathbf{C})| \neq |\phi_j(\mathbf{C})|$.*

In order to establish this lemma, we first prove the following lemma.

Lemma 5.13 *Let $\phi_1(S_1)$ and $\phi_2(S_2)$ be two pp-formulas over a vocabulary τ that are not semi-counting equivalent. Then there is a structure $\mathbf{D}$ such that for every primitive positive formula ϕ over τ we have $|\phi(\mathbf{D})| > 0$ and $|\phi_1(\mathbf{D})| \neq |\phi_2(\mathbf{D})|$.*

Proof. Let $\mathbf{B}$ be any structure on which ϕ_1 and ϕ_2 have a non-zero but different number of solutions. Such a structure exists by definition of semi-counting equivalence. We claim that we can choose $\mathbf{D} = \mathbf{B} + k\mathbf{I}$ for some $k \in \mathbb{N}, k > 0$ where $\mathbf{B} + k\mathbf{I}$ is defined as in the proof of Theorem 5.9. By way of contradiction, assume that $|\phi_1(\mathbf{B} + k\mathbf{I})| = |\phi_2(\mathbf{B} + k\mathbf{I})|$ for all $k \in \mathbb{N}, k > 0$. Then with the same argument as in the proof of Theorem 5.9 we get the contradiction that $|\phi_1(\mathbf{B})| = |\phi_2(\mathbf{B})|$. □

Proof. (Lemma 5.12) We prove this by induction on n; the case $n = 2$ is implied by Lemma 5.13.

When $n > 2$, we first observe that it suffices to prove the result when the ϕ_i are pairwise not semi-counting equivalent, so we assume that this holds. Let $\mathbf{D}$ be the structure that we get by induction for $\phi_1, \ldots, \phi_{n-1}$. We may assume w.l.o.g. that

$$|\phi_1(\mathbf{D})| < |\phi_2(\mathbf{D})| < \ldots < |\phi_{n-1}(\mathbf{D})|.$$

If it holds that $|\phi_n(\mathbf{D})| \neq |\phi_i(\mathbf{D})|$ for every $i \in [n-1]$, then we are done. So we assume that there is an index $i \in [n-1]$ such that $|\phi_n(\mathbf{D})| = |\phi_i(\mathbf{D})|$.

Let $\mathbf{D}'$ be the structure we get by applying Lemma 5.13 to ϕ_n and ϕ_i.

Now choose k such that for every j with $1 < j \leqslant i$ we have

$$\frac{|\phi_{j-1}(\mathbf{D})|^k}{|\phi_j(\mathbf{D})|^k} < \frac{1}{|\mathrm{lib}(\phi_{j-1})|^{|D'|}}.$$

Then we have for every $\ell \geqslant k$ and $1 < j < i$

$$\begin{aligned}|\phi_{j-1}(\mathbf{D}^\ell \times \mathbf{D}')| &= |\phi_{j-1}(\mathbf{D}^\ell)| \cdot |\phi_{j-1}(\mathbf{D}')| \\ &\leqslant |\phi_{j-1}(\mathbf{D}^\ell)| \cdot |\mathrm{lib}(\phi_{j-1})|^{|D'|} \\ &< |\phi_j(\mathbf{D}^\ell)| \\ &\leqslant |\phi_j(\mathbf{D}^\ell)| \cdot |\phi_j(\mathbf{D}')| \\ &= |\phi_j(\mathbf{D}^\ell \times \mathbf{D}')|.\end{aligned}$$

Analogously, we get for every $\ell > k$ that

$$|\phi_{i-1}(\mathbf{D}^\ell \times \mathbf{D}')| < |\phi_n(\mathbf{D}^\ell \times \mathbf{D}')|.$$

Now choose k' such that for every j with $i \leqslant j < n$ we have

$$\frac{|\phi_{j+1}(\mathbf{D})|^{k'}}{|\phi_j(\mathbf{D})|^{k'}} > |\mathrm{lib}(\phi_j)|^{|D'|}.$$

Then we have for every $\ell > k'$ and every $i \leqslant j < n$

$$\begin{aligned}|\phi_j(\mathbf{D}^\ell \times \mathbf{D}')| &= |\phi_j(\mathbf{D}^\ell)| \cdot |\phi_j(\mathbf{D}')| \\ &\leqslant |\phi_j(\mathbf{D}^\ell)| \cdot |\mathrm{lib}(\phi_j)|^{|D'|} \\ &< |\phi_{j+1}(\mathbf{D}^\ell)| \\ &\leqslant |\phi_{j+1}(\mathbf{D}^\ell)| \cdot |\phi_j(\mathbf{D}')| \\ &= |\phi_{j+1}(\mathbf{D}^\ell \times \mathbf{D}')|.\end{aligned}$$

Similarly, we get for every $\ell > k$ that

$$|\phi_{i+1}(\mathbf{D}^\ell \times \mathbf{D}')| > |\phi_n(\mathbf{D}^\ell \times \mathbf{D}')|.$$

Now choosing $\ell = \max(k, k')$ and noting that

$$\begin{aligned}|\phi_i(\mathbf{D}^\ell \times \mathbf{D}')| &= |\phi_i(\mathbf{D}^\ell)| \cdot |\phi_i(\mathbf{D}')| \\ &\neq |\phi_n(\mathbf{D}^\ell)| \cdot |\phi_n(\mathbf{D}')| \\ &= |\phi_n(\mathbf{D}^\ell \times \mathbf{D}')|\end{aligned}$$

completes the proof with $\mathbf{C} = \mathbf{D}^\ell \times \mathbf{D}'$. □

The following is a consequence of this lemma.

Lemma 5.14 *Let $\phi_1(S_1), \ldots, \phi_n(S_n)$ be connected, liberal pp-formulas over the same vocabulary τ that are pairwise not counting equivalent. Then there exists a structure $\mathbf{C}$ (over τ) such that*

- *for all pp-formulas ϕ (over τ) it holds that $|\phi(\mathbf{C})| > 0$, and*
- *for all distinct $i, j \in [n]$, it holds that $|\phi_i(\mathbf{C})| \neq |\phi_j(\mathbf{C})|$.*

Proof. By Lemma 5.12, it suffices to show that the pp-formulas ϕ_i are pairwise not semi-counting equivalent. Since each ϕ_i is connected and liberal, we have $\phi_i = \widehat{\phi}_i$. Thus, by the hypothesis that the ϕ_i are pairwise not counting equivalent in combination with Theorem 5.9, we obtain that the ϕ_i are pairwise not semi-counting equivalent. □

5.3 The all-free case

The aim of this subsection is the proof of Theorem 3.1 in the special case of all-free ep-formulas. Recall that an ep-formula is *all-free* if it is the disjunction of prenex pp-formulas, each of which is *free* in that it has a non-empty set of free variables. We will later in Section 5.4 use the result on all-free formulas to prove the general version of Theorem 3.1.

For every $\phi(V) \in \Phi$ we define a set ϕ^* of free pp-formulas; then, we define $\Phi^* = \bigcup_{\phi \in \Phi} \phi^*(V)$. Let $\phi(V) = \phi_1(V) \vee \ldots \vee \phi_s(V)$ where the $\phi_i(V)$ are free pp-formulas. By inclusion-exclusion we have for every structure $\mathbf{B}$ that

$$\begin{aligned} |\phi(\mathbf{B})| &= \sum_{J \in [s]} (-1)^{|J|+1} |(\bigwedge_{j \in J} \phi_j)(\mathbf{B})| \\ &= \sum_{J \in [s]} (-1)^{|J|+1} |\phi_J(\mathbf{B})|, \end{aligned} \tag{1}$$

where the $\phi_J(V) = \bigwedge_{j \in J} \phi_j(V)$ are pp-formulas. Now iteratively do the following: If there are two summands $c|\phi_J(\mathbf{B})|$ and $c'|\phi_{J'}(\mathbf{B})|$ such that ϕ_J and $\phi_{J'}$ are counting equivalent, delete both summands and add $(c + c')|\phi_J|$ to the sum. When this operation can no longer be applied, delete all summands with coefficient zero. The pp-formulas that remain in the sum form the set ϕ^*.

Example 5.15 It shall be advantageous to again consider Example 4.2. There we started off with

$$\phi(V) = \phi_1(V) \vee \phi_2(V) \vee \phi_3(V).$$

Inclusion-exclusion yields

$$\begin{aligned} |\phi(\mathbf{B})| =& |\phi_1(\mathbf{B})| + |\phi_2(\mathbf{B})| + |\phi_3(\mathbf{B})| \\ & - |(\phi_1 \wedge \phi_2)(\mathbf{B})| - |(\phi_1 \wedge \phi_3)(\mathbf{B})| \\ & - |(\phi_2 \wedge \phi_3)(\mathbf{B})| + |(\phi_1 \wedge \phi_2 \wedge \phi_3)(\mathbf{B})|. \end{aligned}$$

Now we simplify as described above and get

$$|\phi(\mathbf{B})| = 3 \cdot |\phi_1(\mathbf{B})| - 2 \cdot |(\phi_1 \wedge \phi_3)(\mathbf{B})|.$$

Consequently, for this example we have

$$\phi^* = \{\phi_1, \phi_1 \wedge \phi_3\}.$$

□

The algorithm discussed above directly yields the following proposition.

Proposition 5.16 *There exists an algorithm that,*
when an all-free ep-formula ϕ is given as input, outputs a set $\phi^* := \{\phi_1^*, \ldots, \phi_\ell^*\}$ of free pp-formulas, which are pairwise not counting equivalent, and coefficients $c_1, \ldots, c_\ell \in \mathbb{Z}\backslash\{0\}$ such that for every structure $\mathbf{B}$, $|\phi(\mathbf{B})| = \sum_{i=1}^{\ell} c_i |\phi_i^*(\mathbf{B})|$.

We will also require the following two facts for our proof.

Proposition 5.17 *Let us presume that $\phi(S)$ and $\phi'(S')$ are two semi-counting equivalent free pp-formulas that are not counting equivalent and let $(\mathbf{A}, S)$ and $(\mathbf{A}', S')$ be the structures of ϕ and ϕ', respectively. Then $\mathbf{A}$ and $\mathbf{A}'$ are not homomorphically equivalent.*

Proof. $\phi(S)$ and $\phi'(S')$ are semi-counting equivalent, so we have by Theorem 5.9 and Theorem 5.4 that $\widehat{\phi(S)}$ and $\widehat{\phi'(S')}$ are renaming equivalent. It follows that $\mathbf{A}$ and $\mathbf{A}'$ are homomorphically equivalent via homomorphisms $h : A \to A'$, $h' : A' \to A$ that act as bijections between S and S'.

If there exists a homomorphism g from $\mathbf{A}$ to $\mathbf{A}'$, then we can extend h (using the definition of g) to be defined on the components of ϕ deleted in the construction of $\widehat{\phi}$, to obtain a homomorphism from $\mathbf{A}$ to $\mathbf{A}'$ extending h. If there exists a homomorphism g' from $\mathbf{A}'$ to $\mathbf{A}$, we can extend h' in an analogous way. However, the existence of both such extensions would imply by definition that $\phi(S)$ and $\phi'(S')$ are counting equivalent. We may thus conclude that either there is no homomorphism $\mathbf{A} \to \mathbf{A}'$ or there is no homomorphism $\mathbf{A}' \to \mathbf{A}$. □

Lemma 5.18 *There is an oracle FPT-algorithm that performs the following: given a set $\phi_1, \ldots, \phi_s$ of semi-counting equivalent free pp-formulas that are pairwise not counting equivalent, a sequence $c_1, \ldots, c_s \in \mathbb{Z}\backslash\{0\}$, and a structure $\mathbf{B}$, the algorithm computes $|\phi_i(\mathbf{B})|$ for every $i \in [s]$; it may make calls to an oracle that provides $\sum_{i=1}^{s} c_i \cdot |\phi_i(\mathbf{B}')|$ upon being given a structure $\mathbf{B}'$. Here, the ϕ_i with the c_i constitute the parameter.*

To establish this lemma, we first demonstrate the following proposition.

Proposition 5.19 *Let $\phi_1, \ldots, \phi_s$ be a sequence of semi-counting equivalent pp-formulas that are pairwise not counting equivalent. Then there is a structure $\mathbf{C}$ and $i \in [s]$ such that $\mathbf{C} \models \phi_i$ but $\mathbf{C} \not\models \phi_j$ for all $j \in [s]\backslash\{i\}$.*

Proof. Let $\mathbf{A}_1, \ldots, \mathbf{A}_n$ be the structures of the queries $\phi_1, \ldots, \phi_n$. By Proposition 5.17 the structures $\mathbf{A}_i$ are pairwise not homomorphically equivalent. For $i, j \in [n]$, we write $\phi_i < \phi_j$ if there is a homomorphism from $\mathbf{A}_i$ to $\mathbf{A}_j$. It is easy to check that $<$ induces a partial order on the ϕ_i. Let ϕ_i be a minimal element of this partial order, then there is no homomorphism from any $\mathbf{A}_j$ to ϕ_i with $i \neq j$. Setting $\mathbf{C} = \mathbf{A}_i$ completes the proof. □

Proof. (Lemma 5.18) We give and algorithm that recursively computes the $|\phi_i(\mathbf{B})|$ one after the other. So let the parameter and the input be given as in

the statement of the lemma. By Proposition 5.19, there is an $i \in [n]$ and a structure $\mathbf{C}$ such that $\mathbf{C} \models \phi_i$ but $\mathbf{C} \not\models \phi_j$ for all $j \in [s] \backslash \{i\}$. W.l.o.g. assume $i = s$. Then $|\phi_i(\mathbf{B} \times \mathbf{C})| = 0$ for $i < s$. Consequently, we have that the oracle lets us compute $c_s \cdot |\phi_n(\mathbf{B} \times \mathbf{C})| = c_s \cdot |\phi_n(\mathbf{B})| \cdot |\phi_n(\mathbf{C})|$. Computing $|\phi_n(\mathbf{C})|$ by brute force then yields $|\phi_s(\mathbf{B})|$.

Now note that for every structure $\mathbf{B}'$ we can also compute $\sum_{i=1}^{s-1} c_i \cdot |\phi_i(\mathbf{B}')|$ by this approach with one subtraction. So we can apply the algorithm again for $\phi_1, \ldots, \phi_{s-1}$, answering oracle queries for the smaller sum $\sum_{i=1}^{s-1} c_i \cdot |\phi_i(\mathbf{B}')|$ with the help of the oracle for $\sum_{i=1}^{s} c_i \cdot |\phi_i(\mathbf{B}')|$. □

We can now prove Theorem 3.1 for all-free ep-formulas.

Theorem 5.20 *Let Φ be a set of all-free ep-formulas. There exists a set Φ^* of free prenex pp-formulas such that the counting case problems* $\mathsf{count}[\Phi]$ *and* $\mathsf{count}[\Phi^*]$ *are equivalent under counting slice reductions.*

Before giving the technical details of the proof of Theorem 5.20, let us first descibe the ideas. The proof follows the approach presented in the examples of Section 4. In particular, the less straightforward reduction from $\mathsf{count}[\Phi^*]$ to $\mathsf{count}[\Phi]$ proceeds as we did in Example 4.3. Given ϕ and ϕ^*, we can evaluate $|\phi'(\mathbf{B})|$ for $\phi' \in \phi^*$ with an oracle for $|\phi(\mathbf{B} \times \mathbf{C}^\ell)|$ for a suitable structure $\mathbf{C}$ as in that example. The main difference is that, instead of having $\mathbf{C}$ explicitly as in Example 4.3, we here know from Lemma 5.12 that a structure $\mathbf{C}$ exists for which all formulas in ϕ^* have a different number of satisfying assignments. We can then compute $\mathbf{C}$ by brute force as it depends only on ϕ. This then allows to compute $\phi'(\mathbf{B})$ by solving a system of linear equations.

We now give the technical detail of the proof.

Proof. Let us first specify the reduction from $\mathsf{count}[\Phi]$ to $\mathsf{count}[\Phi^*]$, which is quite straightforward. The relation U is the set of pairs (ϕ, ϕ^*) such that ϕ is an all-free ep-formula and ϕ^* is the output of the algorithm of Proposition 5.16 on input ϕ. Obviously, this satisfies the coverage condition. Then the oracle-FPT-algorithm to compute $\phi(\mathbf{B})$ given ϕ, ϕ^* and $\mathbf{B}$ first computes all of the $|\phi_i^*(\mathbf{B})|$ by oracle calls and then uses Proposition 5.16. This completes the reduction.

For the other direction, let $\phi' \in \Phi^*$. We set U to be the set of all pairs $(\phi', \{\phi\})$ such that ϕ is an all-free ep-formula and $\phi' \in \phi^*$. Given ϕ', ϕ and $\mathbf{B}$, we compute $|\phi'(\mathbf{B})| := r(\phi', \{\phi\}, \mathbf{B})$ as follows: Let $\phi_1^*, \ldots, \phi_s^*$ be the equivalence classes of ϕ^* with respect to semi-counting equivalence. Now choose a strucuture $\mathbf{C}$ as in Lemma 5.12. Then for $\psi, \psi' \in \phi^*$ we have $|\psi(\mathbf{C})| \neq |\psi'(\mathbf{C})|$ if ψ and ψ' are from different equivalence classes with respect to semi-counting equivalence, and otherwise $|\psi(\mathbf{C})| = |\psi'(\mathbf{C})| > 0$. Fix for each $j \in [s]$ a formula in ϕ_j^* and call it ψ_j. Moreover, denote by c_ψ the coefficiencent of ψ in Proposition 5.16. Using this notation and Proposition 5.16 we obtain, for every $\ell \in \mathbb{N}$, that

$$|\phi(\mathbf{B} \times \mathbf{C}^\ell)| = \sum_{j=1}^{s} |\psi_j(\mathbf{C})|^\ell \Big(\sum_{\psi \in \phi_j^*} c_\psi |\psi(\mathbf{B})| \Big).$$

Note that this is a linear equation where the coefficients have the form $|\psi_j(\mathbf{C})|^\ell$; these can be computed by brute force. Letting ℓ range from 0 to $s-1$ thus yields a system of linear equations whose coefficient matrix is a Vandermonde matrix. Consequently, with s oracle calls we can compute $\sum_{\psi \in \phi_j^*} c_\psi |\psi(\mathbf{B})|$ for each j. We use Lemma 5.18 to compute $\phi'(\mathbf{B})$. □

5.4 The general case

We now indicate how to prove Theorem 3.1 in its full generality.

We may assume that each ep-formula $\phi \in \Phi$ is normalized. For each ep-formula ϕ, define ϕ_{af} to be the all-free part of ϕ, that is, the disjunction of the ϕ-disjuncts that are free; define Φ_{af} to be $\{\phi_{\mathrm{af}} \mid \phi \in \Phi\}$; and, define ϕ_{af}^- to be the set of formulas in ϕ_{af}^* that do not logically entail a sentence disjunct of ϕ. We define ϕ^+ to be the union of ϕ_{af}^- and the set containing each pp-sentence disjunct of ϕ; and, we define Φ^+ to be $\bigcup_{\phi \in \Phi} \phi^+$.

Example 5.21 Set $V = \{w, x, y, z\}$; we consider the formulas $\phi(V) = \phi_1(V) \vee \phi_2(V) \vee \phi_3(V)$ defined in Example 4.2. Define

$$\theta_1(V) = \exists a \exists b \exists c \exists d E(a,b) \wedge E(b,c) \wedge E(c,d),$$

and define

$$\theta(V) = \phi_1(V) \vee \phi_2(V) \vee \phi_3(V) \vee \theta_1(V).$$

The all-free part of θ is $\theta_{\mathrm{af}} = \phi_1(V) \vee \phi_2(V) \vee \phi_3(V)$, since each of these three disjuncts has a non-empty set of free variables, whereas θ_1 has an empty set of free variables. According to Example 5.15, we have

$$\theta_{\mathrm{af}}^* = \{\phi_1, \phi_1 \wedge \phi_3\}.$$

Now, observe that $\phi_1 \wedge \phi_3$ logically entails the sentence disjunct θ_1 of θ; on the other hand, ϕ_1 does not logically entail θ_1. Hence, we have that $\theta_{\mathrm{af}}^- = \{\phi_1\}$. We have θ^+ to be the union of θ_{af}^- and $\{\theta_1\}$, so

$$\theta^+ = \{\phi_1, \theta_1\}.$$

□

The idea of the proof of Theorem 3.1 is as follows. The counting slice reduction from $\mathsf{count}[\Phi]$ to $\mathsf{count}[\Phi^+]$ has U as the set of pairs (ϕ, ϕ^+) where ϕ is a normalized ep-formula; r on $(\phi(V), \phi^+, \mathbf{B})$ behaves as follows. First, it checks if there is a sentence disjunct θ of ϕ that is true on $\mathbf{B}$; if so, it outputs $|B|^{|V|}$; otherwise, it makes use of the counting slice reduction from $\mathsf{count}[\Phi_{\mathrm{af}}]$ to $\mathsf{count}[\Phi_{\mathrm{af}}^*]$. The counting slice reduction from $\mathsf{count}[\Phi^+]$ to $\mathsf{count}[\Phi]$ has U as the set $\{(\psi, \{\phi\}) \mid \psi \in \phi^+\}$; r on $(\psi(V), \phi(V), \mathbf{B})$ is defined as follows. When $\psi \in \phi_{\mathrm{af}}^-$, the counting slice reduction (U', r') from $\mathsf{count}[\Phi_{\mathrm{af}}^*]$ to $\mathsf{count}[\Phi_{\mathrm{af}}]$ is used to determine $|\psi(\mathbf{B})|$; this is performed by passing to r' a treated version of $\mathbf{B}$, on which no sentence disjunct of ϕ may hold. When ψ is a sentence disjunct of ϕ, an oracle query is made to obtain the count of ϕ on a treated version of $\mathbf{B}$; on this treated version, it is proved that all assignments satisfy ϕ if and only if $\mathbf{B} \models \psi$.

6 Conclusion

We have shown a trichotomy for the parameterized complexity of counting satisfying assignments to existential positive formulas of bounded arity. To this end, the main technical contribution was the equivalence theorem (Theorem 3.1) stating that for every set of existential positive formulas there is a set of primitive positive formulas that is computationally equivalent with respect to the counting problem studied. After showing this equivalence theorem, we could derive our trichotomy in a rather straightforward fashion by invoking a previous trichotomy for primitive positive formulas as a black-box.

In order to prove the equivalence theorem, we gave a syntactic characterization for when two pp-formulas are counting equivalent, that is, have the same number of satisfying assignments with respect to every finite structure. This result can be seen as an adaption, to the counting setting, of classical work of Chandra and Merlin [CM77] that characterizes logical equivalence of primitive positive formulas.

Let us note that the assumption of bounded arity is not needed in the proof of the equivalence theorem. It only appears in our trichotomy theorem because it is already present in the previous trichotomy on primitive positive formulas that we use. Consequently, if one could adapt the work of Marx [Mar10b] on model checking unbounded arity primitive positive formulas to counting to show a dichotomy or trichotomy for counting, this would directly give the corresponding result for existential positive formulas by applying our equivalence theorem.

Finally, let us remark that we are not aware of any fragment of first-order logic extending existential positive queries for which even model checking is understood, from the viewpoint of classifying the complexity of all sets of queries (for more information, see the discussion in the introduction of the article [Che14b]). Hence, the research project of extending our complexity classification beyond existential positive queries would first require an advance in the study of model checking in first-order logic.

References

[AHV95] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.

[CG10] Hubie Chen and Martin Grohe. Constraint satisfaction with succinctly specified relations. *Journal of Computer and System Sciences*, 76(8):847–860, 2010.

[Che14a] Hubie Chen. On the complexity of existential positive queries. *ACM Trans. Comput. Log.*, 15(1), 2014.

[Che14b] Hubie Chen. The tractability frontier of graph-like first-order query sets. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, page 31, 2014.

[CM77] Ashok K. Chandra and Philip M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceddings of STOC'77*, pages 77–90, 1977.

[CM14a] Hubie Chen and Stefan Mengel. A trichotomy in the complexity of counting answers to conjunctive queries. *CoRR*, abs/1408.0890, 2014.

[CM14b] Hubie Chen and Moritz Müller. One hierarchy spawns another: graph deconstructions and the complexity classification of conjunctive queries. In *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS '14, Vienna, Austria, July 14 - 18, 2014*, pages 32:1–32:10, 2014.

[CM15] Hubie Chen and Stefan Mengel. A trichotomy in the complexity of counting answers to conjunctive queries. In *18th International Conference on Database Theory, ICDT 2015, March 23-27, 2015, Brussels, Belgium*, pages 110–126, 2015.

[DJ04] Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004.

[DM13] A. Durand and S. Mengel. Structural tractability of counting of solutions to conjunctive queries. In *Proceedings of the 16th International Conference on Database Theory (ICDT 2013)*, 2013.

[FG06] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.

[Gro07] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *Journal of the ACM*, 54(1), 2007.

[GS14] Gianluigi Greco and Francesco Scarcello. Counting solutions to conjunctive queries: structural and hybrid tractability. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 132–143, 2014.

[GSS01] Martin Grohe, Thomas Schwentick, and Luc Segoufin. When is the evaluation of conjunctive queries tractable? In *STOC 2001*, 2001.

[Mar10a] Dániel Marx. Can you beat treewidth? *Theory of Computing*, 6(1):85–112, 2010.

[Mar10b] Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 735–744, 2010.

[PS11] Reinhard Pichler and Sebastian Skritek. Tractable counting of the answers to conjunctive queries. In *Proceedings of the 5th Alberto Mendelzon International Workshop on Foundations of Data Management*, 2011.

[PY99] C. Papadimitriou and M. Yannakakis. On the Complexity of Database Queries. *Journal of Computer and System Sciences*, 58(3):407–427, 1999.

A Proof of Theorem 3.1

Proof. (Theorem 3.1) We first describe a counting slice reduction (U, r) from $\mathsf{count}[\Phi]$ to $\mathsf{count}[\Phi^+]$. Let (U', r') be the counting slice reduction from $\mathsf{count}[\Phi_{\text{af}}]$ to $\mathsf{count}[\Phi^*_{\text{af}}]$ given by Theorem 5.20. Define U to be the set $\{(\phi, \phi^+) \mid \phi \text{ is a normalized ep-formula }\}$. When $(\phi, \phi^+) \in U$, we define $r(\phi(V), \phi^+, \mathbf{B})$ to be the result of the following algorithm, which is FPT with respect to (π_1, π_2). For each sentence disjunct θ of $\phi(V)$, the algorithm queries $\mathsf{count}(\theta, \mathbf{B})$; if for some such disjunct θ it holds that $\mathbf{B} \models \theta$, then the algorithm outputs $|V|^{|B|}$. Otherwise, for any assignment $f : V \to B$, it holds that $\mathbf{B}, f \models \phi$ if and only if $\mathbf{B}, f \models \phi_{\text{af}}$. So, the algorithm returns $r'(\phi_{\text{af}}, \phi^*_{\text{af}}, \mathbf{B})$ by running the corresponding algorithm for r'. In this run, the algorithm for r' only makes queries of the form $(\psi, \mathbf{B})$ (with $\psi \in \phi^*_{\text{af}}$); such queries where $\psi \in \phi^-_{\text{af}}$ are resolved using the oracle in the definition of counting slice reduction, and queries where $\psi \in \phi^*_{\text{af}} \backslash \phi^-_{\text{af}}$ are answered with 0. Correctness is straightforward to verify.

We next describe a counting slice reduction (U, r) from $\mathsf{count}[\Phi^+]$ to $\mathsf{count}[\Phi]$. Let (U', r') denote the counting slice reduction from $\mathsf{count}[\Phi^*_{\text{af}}]$ to $\mathsf{count}[\Phi_{\text{af}}]$ given by Theorem 5.20. We define $U := \{(\psi, \{\phi\}) \mid \psi \in \phi^+\}$. We need to define $r(\psi(V), \phi(V), \mathbf{B})$ when $(\psi, \phi) \in U$.

Let us describe first an algorithm for the mapping r in the case that $\psi \in \phi^-_{\text{af}}$. Let $(\mathbf{C}_1, V), \ldots, (\mathbf{C}_m, V)$ denote the pp-formulas in ϕ^-_{af}, and let $\mathbf{C}$ denote the disjoint union of the structures $\mathbf{C}_i$. Observe that for any structure $\mathbf{D}$, it holds that $\mathbf{D} \times \mathbf{C}, f \models \phi$ if and only if $\mathbf{D} \times \mathbf{C}, f \models \phi_{\text{af}}$, since no sentence disjunct of ϕ holds on $\mathbf{C}$ (due to the definitions of $\mathbf{C}$ and ϕ^-_{af}). Call the algorithm for r' to compute $r'(\psi, \{\phi_{\text{af}}\}, \mathbf{B} \times \mathbf{C}) = |\psi(\mathbf{B} \times \mathbf{C})|$; note that the oracle queries made by this algorithm can be resolved by an oracle for $\mathsf{count}(\phi, \cdot)$, since all such oracle queries have the form $\mathsf{count}(\phi_{\text{af}}, \cdot \times \mathbf{C})$. As $|\psi(\mathbf{B} \times \mathbf{C})| = |\psi(\mathbf{B})| \cdot |\psi(\mathbf{C})|$, by dividing this quantity by $|\psi(\mathbf{C})|$, one can determine $|\psi(\mathbf{B})|$, which is the desired value. Note that by the definition of $\mathbf{C}$, it holds that $|\psi(\mathbf{C})|$ is non-zero.

In order to describe the behavior of the algorithm for r in the case that ψ is a sentence disjunct of ϕ, we establish the following claim. Let $(\mathbf{A}, V)$ be the structure view of ψ.

Claim: Let $i : V \to V$ be the identity map on V. For each disjunct θ of ϕ, it holds that $\mathbf{A}, i \models \theta(V)$ if and only if $\theta = \psi$.

The backwards direction is clear, so we prove the forwards direction. If a disjunct θ is a free pp-formula, then $\mathbf{A}, i \not\models \theta(V)$ since θ contains an atom using a variable $v \in V$, whereas no tuple of a relation of $\mathbf{A}$ contains any variable from V. If a disjunct θ is a pp-sentence $(\mathbf{A}', V)$ not equal to ψ, then by definition of *normalized* ep-formula, there is no homomorphism from $\mathbf{A}'$ to $\mathbf{A}$ and hence $\mathbf{A}, i \not\models \theta(V)$. This establishes the claim.

Now suppose that ψ is a sentence disjunct of ϕ. In this case, the algorithm for $r(\psi(V), \phi(V), \mathbf{B})$ behaves as follows. It queries $\mathsf{count}(\phi, \mathbf{A} \times \mathbf{B})$ to determine $|\phi(\mathbf{A} \times \mathbf{B})|$; it outputs $|B|^{|V|}$ if $|\phi(\mathbf{A} \times \mathbf{B})|$ is equal to $(|A| \cdot |B|)^{|V|}$ (the *maximum count* possible there), and outputs 0 otherwise. We prove that this is correct by showing that $|\phi(\mathbf{A} \times \mathbf{B})|$ is the maximum count if and only if $\mathbf{B} \models \psi$.

For the backwards direction, suppose that $\mathbf{B} \models \psi$, and denote ψ by $(\mathbf{A}, V)$.

Then, there is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$, and hence there is a homomorphism from $\mathbf{A}$ to $\mathbf{A} \times \mathbf{B}$. It follows that for any assignment $f : V \to V$, one has $\mathbf{A} \times \mathbf{B}, f \models \psi(V)$. For the forwards direction, suppose that $|\phi(\mathbf{A} \times \mathbf{B})|$ is the maximum count. Let $i' : V \to A \times B$ be any map such that for each $v \in V$, the value $i'(v)$ has the form $(i(v), j(v))$ where $j : V \to B$ is a map. We have that $\mathbf{A} \times \mathbf{B}, i' \models \phi(V)$. It follows that there is a disjunct θ of ϕ such that $\mathbf{A} \times \mathbf{B}, i' \models \theta(V)$. It follows that $\mathbf{A}, i \models \theta(V)$ and $\mathbf{B}, j \models \theta(V)$. By the claim established above, we have that $\theta = \psi$. Then, it holds that $\mathbf{B}, j \models \psi$, and we are done. $\square$